company essentially must either continually "re-commit" its in-service base of special access circuits or incur significant cost increases. If a company wanted to move a significant volume of circuits to an alternate provider, that carrier would be unable to do so without substantial penalties unless the migrations were performed "overnight" at the end of its term of agreement in order to maintain compliance with the minimum commitment thresholds (*i.e.* 90%) contained in the tariff contracts.[16] If a customer signed a new contract and then migrated substantial services to an alternate provider, it would incur substantial penalties to reduce its commitment threshold, as discussed above.

Most DS1 service agreements are structured as overall volume commitments for these services, providing "portability" that permits a customer to move or add or disconnect a circuit covered by the agreement without penalty as long as the aggregate active circuit count does not drop below the commitment minimum. (Volume commitments are established on a geographical basis and may be set at the state level or a regional level.) There are exceptions, however. AT&T's discount plans in some of its regions are on an individual circuit basis, and therefore there is no portability. AT&T PacBell offers separate "portability" and individual circuit discount plans which can be entered into concurrently; this combination of contracts achieves the relevant portability and discount conditions. However, in addition to a shortfall penalty if the circuit volume drops below 80% of the minimum committed volume, AT&T PacBell also charges an

[16] While "overnight migration" to an alternative provider is not a realistic practice on a large scale, this approach can be utilized on a limited basis at the end of terms of agreement to optimize a carrier's cost structure by maximizing savings and minimizing penalties.

Excess Circuit penalty. This penalty is equal to the non-recurring channel termination charge ($900), per circuit, and applies to new special access circuit additions that exceed 124% of the committed volume, forcing a customer to increase its minimum commitment or incur substantial liability. The majority of the other incumbent LECs either (1) provide the discount on the committed volume (*e.g.*, BellSouth), or (2) discount the actual volume but require an increased commitment within a set period (*e.g.*, AT&T, Ameritech, and SBC). AT&T PacBell charges an additional penalty if a customer adds new special access circuits that exceed a specified limit without an increased commitment in the portability plan.

At a DS3 level, portability is provided by only a few vendors depending on how the tariffs are structured. AT&T does not provide DS3 portability.

Finally, incumbent LECs may enter into special contract tariffs with customers if those customers can meet certain highly specified criteria. The terms of these specific contracts are unique, and they are generally tailored to a single customer. Based on XO's experience, these contracts are extremely difficult to negotiate. Customers must commit to varying spending levels and changes in operational relationships to gain small incremental discounts beyond the standard term plan discounts.

d. What discounts from tariff "rack rates" of list prices are available for other services such as OCn, Fast Ethernet, or Gigabit Ethernet? Does the availability of discounts vary by geography or density zone, if so, by how much? Do these discounts vary when competitive alternatives are present, and if so, by how much?

Table 1 contains examples of prices for OC-3 and OC-12 services purchased from incumbent LECs across a variety of geographic areas, as well as prices for Ethernet 100

Megabit and Ethernet 1 Gig services. Table 1 shows the percentage differences between month-to-month rates (*i.e.* "rack rates") and the rates for these services under available term plan or contracts. As Table 1 demonstrates, the discounts received for purchasing capacity under a term plan rather than on a month-to-month basis are substantial. Fixed mileage discounts for OC-3 services range from **[Begin Confidential]** [redacted], **[End Confidential]** with variable mileage discounts for OC-3 ranging from **[Begin Confidential]** [redacted] **[End Confidential]**. For OC-12 services, fixed mileage discounts range from **[Begin Confidential]** [redacted] **[End Confidential]**, and variable mileage discounts range from **[Begin Confidential]** [redacted] **[End Confidential]**. For Ethernet 100 Mb service, there is a channel term discount of **[Begin Confidential]** [redacted] **[End Confidential]**, while for Ethernet 1 Gig service the discount is **[Begin Confidential]** [redacted] **[End Confidential]**.

While the availability of discounts does not currently vary by geography or density zone, they do vary when competitive alternatives are present. Individual Case Basis ("ICB") prices are available from the incumbent LECs subject to negotiation and customers' acceptance of certain specific terms and conditions. If presented with information regarding competitive bids, incumbent LECs sometimes offer discounts and/or special promotions to match those competitors.

e. **Do broadband Internet access service providers purchase circuit-mode services for the purposes of second mile and/or middle mile transport? If so, in what circumstances do they do so and to what extent? What are the costs associated with converting a DS1, DS3, or OCn circuit to an IP connection? How much of these costs would be avoided if a broadband ISP had the ability to purchase a transmission service that obviated the need for such conversion, such as Fast Ethernet, Gigabit Ethernet, or**

other packet-switched service for both the second and middle mile transport? Is functionality lost during these multiple conversions?

Broadband Internet access service providers purchase circuit-mode services either (1) to reach end-user customer premises that are not served by their own network facilities, or (2) where a customer premises is served by an incumbent LEC end office that does not have collocation equipment installed by the purchaser or any non-incumbent LEC. **[Begin Confidential]** [redacted] **[End Confidential]** Additionally, XO's network is a SONET/TDM backbone network which transports packets via Sonet. This type of service has historically been purchased because, as XO describes above, Ethernet services either have not been commercially available ubiquitously or, as described above, have not met XO's and other competitive carriers' technical and operational requirements.

With respect to the conversion of DS1, DS3, or OCn circuits, there are many vendors that supply equipment that converts DS1, DS3, OCn signals to Ethernet. There is no loss of functionality in conjunction with such conversion. The general technology is commonly referred to as EoS (Ethernet Over Serial). The specific costs for the conversion of these circuits are dictated by the engineering and equipment requirements of the particular scenario.

XO currently uses the equipment vendor **[Begin Confidential]** [redacted] **[End Confidential]** to perform conversions of DS1, DS3, and OCn circuits to Ethernet. **[Begin Confidential]** [redacted]

REDACTED –
FOR PUBLIC INSPECTION

[redacted]

[End Confidential]

f. Given current and projected demand and supply conditions, what portion of the overall cost of providing broadband Internet service to an end user is attributable to middle mile and second mile transport? Does this portion of cost vary by distance or length of the circuit, and to what extent?

To assess the portion of overall costs that are attributable to middle mile and second mile transport (as defined in the *Public Notice*), XO examined the costs associated with its provision of broadband Internet access service to retail customers over DS1 facilities.[17] XO analyzed the cost of providing service to certain XO customers (1) located in close proximity to XO's network (*i.e.*, its collocation facilities at an incumbent LEC central office) and (2) located further from XO's network and therefore classified as being in a "red zone" outside XO's service area. (Internally, XO refers to such customers as "in LATA" or "out of LATA".) XO also calculated its non-transport

[17] As indicated at note 5 *supra*, XO in this analysis assigns inter-office transport channel mileage costs to the second mile category.

costs associated with this service, including loop costs, direct capital costs, sales and marketing costs, and general administrative costs.

As shown in Tables 3 and 4 (attached in Appendix B), second mile and middle mile transport costs represent a substantial proportion of XO's overall cost of providing service to XO's "out of LATA" customers.[18] Table 3 shows the cost data associated with XO's provision of **[Begin Confidential]** [redacted] **[End Confidential]** service to **[Begin Confidential]** [redacted] **[End Confidential]** customer in Minnesota where the middle mile segment is between **[Begin Confidential]** [redacted] **[End Confidential]** miles in length. In that example, **[Begin Confidential]** [redacted] **[End Confidential]** of XO's overall costs are attributable to middle mile transport and **[Begin Confidential]** [redacted] **[End Confidential]** of its overall costs are attributable to second mile transport.[19] This means that almost **[Begin Confidential]** [redacted] **[End Confidential]** of XO's costs for that service are attributable to its purchase of transport. Not surprisingly, as the middle mile segment becomes longer, the portion of XO's costs attributable to transport also increases. As shown in Table 4, for service to another customer involving a middle mile segment between **[Begin Confidential]** [redacted] **[End Confidential]** miles in length,

---

[18] The data in Table 4 is derived from specific examples for each of the customer categories shown in the tables (*e.g.*, less than 50 miles related to the middle mile, between 50 miles and 100 miles related to the middle mile, etc.) Table 3 provides expanded data for one of these examples.

[19] In this example, XO and its customer have a **[Begin Confidential]** [redacted] **[End Confidential]** service agreement. Typically, the percentage of overall cost attributable to middle mile and second mile transport increases slightly with the length of the customer contract, since XO can spread other upfront expenditures over those longer time frames.

middle mile transport costs constitute **[Begin Confidential]** [redacted] **[End Confidential]** of XO's overall costs. For service to this customer, middle mile and second mile transport together represent almost **[Begin Confidential]** [redacted] **[End Confidential]** of XO's overall costs.

4. Economics of Deployment.

> **c. What are the categories of the capital expenses of constructing second mile and/or middle mile transport? What are the categories of operating expenses of operating second mile and/or middle mile transport? On a per-mile basis, what are the levels of each of those categories of capital and operational expenses? What are the primary factors that affect these costs? How does distance from the nearest Internet point of presence affect each of these categories of costs? Which, and what percent, of these costs are affected by rights of way and pole attachment fees and charges? Do these charges vary based upon who owns or controls the pole, duct, conduit, or right of way? Do these costs vary by percentage of outside plant that is aerial, buried, or conduit, and if so, by how much?**

The analysis that XO recently performed in assessing alternative approaches to expanding service in Charlotte, North Carolina provides a useful illustration of how one carrier compares the relative advantages of using incumbent LEC special access services to meet middle mile service needs with other available alternatives. XO's new Charlotte network will consist of a series of transport rings (totaling **[Begin Confidential]** [redacted] **[End Confidential]** route miles) with varying amounts of fiber on which are located XO POPs, carrier hotels, and incumbent LEC central offices.[20] Prior to initiating deployment of this network, XO examined the costs associated with three different build-out options: (1) self-construction of the facilities; (2) acquisition of dark fiber facilities from a

[20] XO expects that this fiber network will become operational in **[Begin Confidential]** [redacted]. **[End Confidential]**

competitive fiber provider and attaching its own electronics; and (3) using the incumbent LEC's special access services. XO's analysis indicated that relying on the competitive provider's dark fiber facilities would be the least expensive and most cost-effective option. Below, XO reviews the capital and operating numbers that drove this analysis.

Self-Build. Based on its experience in other projects, XO relied on a rough cost estimate of **[Begin Confidential]** [redacted] **[End Confidential]** per foot for a Charlotte, NC self-build. Under this model, the total build-out cost would be approximately **[Begin Confidential]** [redacted] **[End Confidential]** million. Such a large investment would tie up XO's capital, and XO would incur significant carrying costs to finance the project. While following construction XO could in theory resell the excess capacity to other carriers, XO viewed this approach as highly risky given the uncertain nature of bandwidth demands in the Charlotte market.

Competitive Provider's Dark Fiber. A competitive fiber provider offered XO dark fiber facilities for the construction of seven rings. The number of fiber strands in the rings varied from **[Begin Confidential]** [redacted] **[End Confidential]** (totaling **[Begin Confidential]** [redacted] **[End Confidential]** fiber strands), and the route mileage ranged from **[Begin Confidential]** [redacted] **[End Confidential]** miles. The competitive fiber provider proposed that XO enter into a **[Begin Confidential]** [redacted] **[End Confidential]** XO determined that the capital investment required for the electronics to light these facilities and related items would amount to an additional **[Begin Confidential]** [redacted] **[End Confidential]** million, resulting in a total investment of well over **[Begin Confidential]** [redacted] **[End Confidential]** million. Further, XO could readily expand and upgrade these facilities. Having access to fiber at reasonable rates would allow XO to deploy additional SONET rings at an estimated investment cost of **[Begin Confidential]** [redacted] **[End Confidential]** per node. XO could also deploy underlying Dense Wavelength Division Multiplexing ("DWDM") technology as needed to facilitate access to multiples of O-C48 or OC-192 worth of bandwidth using a pair of fibers. XO estimated these annual operating expenses for the competitive fiber provider's facilities as:

1. **[Begin Confidential]** [redacted] **[End Confidential]** annually paid to the competitive fiber provider for fiber maintenance fees.
2. Assuming there are ten nodes on the network, annual payments of approximately **[Begin Confidential]** [redacted] **[End Confidential]** to the incumbent LEC for collocation rent, power, and other fees.

Total annual operating expenses under this option would exceed **[Begin Confidential]** [redacted], **[End Confidential]** which, as described below, is less than **[Begin Confidential]** [redacted] **[End Confidential]** expenditures on incumbent LEC special access services for the same network configuration.

Incumbent LEC Special Access. To assess the price of using incumbent LEC special access facilities, XO examined the prices for **[Begin Confidential]** [redacted]

REDACTED –
FOR PUBLIC INSPECTION

[redacted]

[End Confidential] On this basis, XO determined that reliance on the incumbent LEC special access services in this instance would be cost-prohibitive. XO has come to a similar conclusion regarding special access services in numerous other markets.

***

XO is encouraged by the competitive fiber provider's presence, but it notes that in most geographic areas there is no analogous provider to offer a genuine alternative to the incumbent LEC. Transport competitors like this provider face substantial obstacles to entry and growth. Capital is often difficult to access at rates that can produce a sufficient return on investment. Rights of way may be costly to access, and incumbent LEC special access terms and conditions often lock in potential purchasers so they cannot move to alternative facilities.

d. Do existing long-haul fiber optic service providers offer either middle mile or second mile transport service to all communities that are passed by their long-haul fiber? Why or why not? What are the cost and economics of building a local "on-ramp" or fiber access point at these locations?

Primarily for cost reasons, long-haul fiber optics service providers today do not typically build "add/drop capability" and provide transport service to all of the intermediate population centers on the path of a route. The deployment of long-haul fiber facilities typically requires the installation of optical amplifiers approximately every 100 km and signal regeneration equipment approximately every 400-600 km. (More advanced regeneration equipment requires that carriers regenerate signals every 1500 km or more.) In order to provide add/drop capability to a town along a long-haul fiber route, however, signal regeneration equipment must be deployed at that location for that specific purpose, no matter how far that location is from other deployed electrical regenerators. Currently, deployment of signal regeneration equipment is considerably more expensive and complex than optical amplification; while optical amplification can be accomplished with equipment taking up one shelf in a rack, the regeneration of all electrical signals on one pair of long-haul fiber can require up to four racks of equipment or more. The electronics required for signal regeneration alone can require an investment of up to $250,000, not including the recurring operating costs of space, power, and cooling in the long-haul "hut" that contains the signal regeneration equipment.

In addition, access to local networks is not always available for long-haul fiber providers. In some cases, there are no local fiber networks within reasonable proximity of the long-haul hut. In those instances, to make use of the add/drop capability, long-haul providers need to (1) build fiber out into the local market, (2) lease network capacity from another provider, or (3) simply sell the middle mile service to the local provider, who builds to the long-haul hut. Any of these options would add significantly to the long-haul provider's costs.

III. CONCLUSION

Following its collection and analysis of the requested information about backhaul, the Commission should move forward expeditiously to promote broadband competition, deployment, and availability throughout the United States.

Respectfully submitted,

Heather Burnett Gold/SJB

Heather Burnett Gold
Senior Vice President
heather.b.gold@xo.com

Lisa R. Youngers
Vice President, Federal Affairs
lisa.r.youngers@xo.com

XO Communications, LLC
13865 Sunrise Valley Drive
Herndon, VA 20171
703-547-2000

November 4, 2009

**REDACTED –
FOR PUBLIC INSPECTION**

Appendix A

REDACTED –
FOR PUBLIC INSPECTION

Before the
FEDERAL COMMUNICATIONS COMMISSION
Washington, DC 20554

In the Matter of	)	
	)	
Impact of Middle and Second Mile Access on Broadband Availability and Deployment: NBP Public Notice #11	)	GN Docket No. 09-47
	)	GN Docket No. 09-51
A National Broadband Plan for Our Future	)	
	)	GN Docket No. 09-137
Inquiry Concerning the Deployment of Advanced Telecommunications Capability to All Americans in a Reasonable and Timely Fashion, and Possible Steps to Accelerate Such Deployment Pursuant to Section 706 of the Telecommunications Act of 1996, as Amended by the Broadband Data Improvement Act	)	

DECLARATION OF RANDY NICKLAS

1. I am the Chief Technology Officer for XO Communications, LLC.

2. My business address is 13865 Sunrise Valley Drive, Herndon, VA 20171.

I. PROJECTED BANDWIDTH REQUIREMENTS FOR BROADBAND CONNECTIVITY

3. The amount of IP bandwidth needed to provide adequate broadband Internet access to the end user depends greatly on the nature of the end users and how those users utilize their Internet access. It is also important to note that what is considered adequate access bandwidth is a function of time and can be difficult to predict. At 1.544 Mbps, T1s today are a common but not particularly fast Internet access circuit used by many small to medium U.S. businesses. In 1988, however, the core of the Internet, the NSF-funded NSFNET backbone service, consisted of single T1 links between the core routers.

Many other more recent examples exist of the Internet backbone speeds of yesterday becoming the Internet access speeds of today. For example, 10 Gbps Internet access is not uncommon for large U.S. corporations now, but circuits of this size were first introduced into commercial Internet backbones in 2000.

A. End User Bandwidth Requirements

4. In order to build up a picture of middle mile capacity requirements, we have to estimate the Internet bandwidth requirements of large classes of users today and in the future – over the next five to ten years. The first class of Internet access users addressed here is residential users. For the broad range of residential users, in addition to the standard applications of web surfing, e-mail, VPN access to corporate intranets, and VoIP-based telephony, there is increasing use of the Internet to deliver video content. This content ranges from low-resolution amateur productions such as those hosted by YouTube with modest bandwidth requirements, to high-definition ("HD") professionally produced productions that may require 8-20 Mbps per video stream or channel, depending on the encoding technique. These video streams can last up to the viewing length of the content if they are not streamed at faster than viewing rate, so a large video stream can last for several hours. This lack of traffic "burstiness" is detrimental to the oversubscription of second and middle mile links.

5. Assuming 20 Mbps per video channel (future encoding technique efficiencies can be presumed to be balanced by higher quality video programming in the future) and the use of three simultaneous channels per residence at any given time, there is a requirement of 60 Mbps of average peak Internet bandwidth demand per residential subscriber for

video applications. To account for non-video applications and provide margin for future growth, a raw residential broadband requirement per subscriber of 100 Mbps is assumed. For residential use, the access bandwidth need not be symmetric, since the video traffic that is the largest component typically flows from the network towards the subscriber, or downstream. Thus, the 100 Mbps requirement is in the downstream direction. Less bandwidth is likely to be necessary in the upstream direction. To adequately support current and future applications (such as video surveillance camera streams or teleconferencing applications), it is likely that an upstream capacity of 10 Mbps will be sufficient, with greater bandwidth capabilities always desirable since, as a general rule, applications appear to saturate currently available bandwidth.

6. The Internet access bandwidth requirements of a business depend on a combination of the nature of the particular business' use of the Internet and the number of employees served by the Internet connection. The principal bandwidth driver for residential Internet access, the delivery of high-quality video programming, cannot be expected to be present to any large degree in the workplace. Even with video conferencing on the rise, the workday of most employees are unlikely to be taken up by teleconferences. Moreover, the bandwidth requirements of an individual in the workplace are currently bursty in nature, and it is reasonable to expect to them to remain so. For example, typical business Internet usage might entail a web search followed by several document downloads and then a relatively lengthy review of the downloaded documents. Besides a relatively modest bandwidth allocation to video programming today and going forward, the only other commonplace streaming Internet application is

VoIP-based telephony, which today requires no more than roughly 100 Kbps per conversation. Combining the less-than-residential bandwidth requirements of the workplace with the typical bursty nature of Internet usage per employee in the workplace, I believe that *an average 0.5-1.0 Mbps of Internet access bandwidth per employee will be adequate for some time to come.* Thus, a 20 employee worksite might be well served with a 10 Mbps Internet access service, while a more bandwidth-intensive organization with 100 employees might need 100 Mbps or more of Internet access bandwidth.

7. The U.S. Census Bureau 2004 data on employment size of US firms indicates an average of roughly 50 employees per business location across all firms with 100 or more employees (*see* http://www.census.gov/epcd/www/smallbus.html for the data underlying this statement). Assuming each of these business locations is independently connected to the Internet, and using the range of Internet access bandwidth per employee described above, one arrives at the estimate of 25-50 Mbps of Internet access bandwidth required per workplace. Of course, many larger enterprises channel their employees' Internet access through fewer and larger access circuits using their own internal IP networks or intranets, and so the figure of 25-50 Mbps of Internet access per enterprise connection is low for these enterprises. To summarize, I believe a lower bound on the Internet access bandwidth requirements for a single business connection going forward is 25 Mbps, and, to account for future growth and uncertainty, a value of *50 Mbps for business Internet access circuit* is used in the following analysis.

B. Middle Mile and Second Mile Bandwidth Requirements

8. These bandwidth estimates above reflect the raw bandwidth delivered to the customer premises by the end-to-end distribution chain from the customer premises to an ISP PE router, modeled by the concatenation of a last mile segment, a second mile segment, and a middle mile segment as defined by the FCC in its *Public Notice*. Turning to the capacity requirements of a middle mile segment, it is important to note that the capacity of such a segment is not necessarily simply the sum of the bandwidth requirements of the individual residential and business customers served by middle mile segment. Since (i) Internet traffic is packetized, (ii) in general the instantaneous bandwidth demand of any customer is independent of the demand of any other customer, and (iii) in many cases individual customer demand is bursty in nature, the peak aggregate bandwidth requirement of all the customers served through a middle mile segment is less than the simple sum of individual peak bandwidth demands. This type of communication link sharing is known as statistical multiplexing, and can result in significant oversubscription of shared link bandwidth when compared to the traditional fixed partitioning of shared link bandwidth used in time division multiplexing ("TDM") networks such as the public switched telephone network ("PSTN"). The Internet backbones and packet-based access networks operated by ISPs today all employ varying degrees of oversubscription, often implemented in several locations or stages in a given IP network architecture, in order to achieve the necessary economics. Before discussing these oversubscription ratios further, it is important to note that (i) many Internet access circuits today are TDM-based and not oversubscribed, and (ii) traffic due to independent

IP video or voice streams cannot be statistically multiplexed, since they are examples of constant or near-constant bit rate traffic – the opposite of traditionally bursty Internet traffic. Thus, a T1 access loop from a customer premises to an ISP PE router admits no oversubscription of its bandwidth allocation even though it carries pure IP traffic. Also, the oversubscription ratios employed to deliver today's Internet traffic mix will not necessarily apply to future traffic mixes that would have a large video streaming component – residential Internet service for example. Further complicating this discussion of statistical multiplexing is the technique of caching popular web content as close to the consumer as possible, as well as the use of multicast distribution of video streams. Given these factors, it can be difficult to make general statements regarding oversubscription for today's Internet access networks and particularly for future access networks.

9. In order to achieve any degree of oversubscription of an access network connecting Internet customers to an ISP's PE router, at least one stage of oversubscription must exist on the last mile + second mile + middle mile chain. To achieve this, a packet switch (Ethernet, MPLS, or IP based) must be placed at the junction of the last mile and second mile segments, at the junction of the second mile and middle mile segments, or at both junction points. It is worth noting that the ISP PE router at the Internet gateway is a packet switch and hence provides a concentration function, allowing for the oversubscription of its uplinks today. This is the first stage of oversubscription in many existing Internet networks.

10. I believe that the network architectures of the future will include packet switches, based on some combination of Ethernet and MPLS, deployed at least to the junction points of the second mile and middle mile access segments: that is, in incumbent LEC central offices, cable network head ends, and mobile switching centers ("MSCs"). It is this deployment of layer 2 packet switches that will allow concentration of Internet traffic and the statistical multiplexing of the aggregate traffic onto the middle mile segment. My estimate of the range of oversubscription ratios that could obtain at this beginning of the middle mile is 4:1 to 10:1, depending on the nature of the customers and the fraction of video in the traffic aggregate. Furthermore, I estimate that each middle mile segment would aggregate 10 to 25 second mile segments, with each second mile segment serving roughly 100 last mile segments in turn.

11. Thus, the number of last mile segments aggregating into a middle mile segment is: (100)*(10 – 25) = 1,000 – 2,500. As described in previous paragraphs, the forward-looking peak bandwidth requirements of the unshared last mile segment are 50 Mbps for business locations and 100 Mbps for residential locations. If all of the last mile segments aggregating into a single middle mile segment serve business locations, then, *assuming no oversubscription*, the required middle mile segment capacity is 1,000 – 2,500)*(50 Mbps) = 50,000 – 125,000 Mbps = 50 – 125 Gbps. Similarly, if all last mile segments serve residential customers, then 100 – 250 Gbps of middle mile segment capacity is required, again under the assumption of no oversubscription. Since the mix of residential and business customers served by one middle mile segment depends on a number of factors, beginning with the population density of the area served by the aggregating

REDACTED –
FOR PUBLIC INSPECTION

central office/cable head end/MSC, it is assumed that each such location has a raw aggregate bandwidth ranging from 50 to 250 Gbps. Finally, taking this interval of raw aggregate bandwidth and applying the range of oversubscription ratios described above (4:1 to 10:1), the aggregate bandwidth required of a middle mile segment ranges from 5 Gbps all the way up to 62.5 Gbps. While some observers may question some of the assumptions underlying this analysis, it appears clear that *the forward-looking aggregate Internet access bandwidth required of a middle mile segment is measured in gigabits per second.*

12. Given the estimate that the middle mile on average aggregates 10 to 25 second mile segments, the bandwidth requirement for the second mile is obviously less than that for the middle mile. In the middle mile model described above, a range of *cumulative* oversubscription ratios (up to and including the junction switch joining second mile and middle mile segments) of 4:1 to 10:1 was postulated. The junction device joining last mile segments into a single second mile segment, resident in a remote terminal enclosure, a cable coax-fiber junction box, or a cell tower base transceiver station (BTS), can in principle be a packet switch device that concentrates packet flows from the last mile segments into the second mile uplink. This device is capable of providing the first stage of statistical multiplexing in the last mile-second mile-middle mile path between the customer and the ISP PE router. It is assumed that half of the end-to-end oversubscription assumed in the middle mile model is obtained in the packet switch concentrator that aggregates last mile segments into a second mile segment. Thus, this first stage of oversubscription ranges from 2:1 to 5:1.

REDACTED –
FOR PUBLIC INSPECTION

13. As indicated above, I estimate that each second mile segment aggregates on average 100 last mile customers. Given the above-described estimates that business customers require 50 Mbps of peak Internet bandwidth while residential customers require 100 Mbps of bandwidth, the offered last mile load ranges from 5 Gbps to 10 Gbps, depending on the mix of residential and business customers. Taking into account the first stage of oversubscription, *the capacity of the second mile ranges from 1 Gbps to 5 Gbps.*

14. These residential and business customer bandwidth requirement estimates may require some modification in the case of wireless last mile applications. While next generation mobile technologies such as WiMax and LTE should be capable of eventually supporting 50-100 Mbps of bandwidth per subscriber in at least the downstream direction, they will require large per channel allocations of 20 MHz or more of mobile spectrum to do so. Moreover, wireless first mile technologies will have difficulty economically supporting streaming video applications such as HD video for large numbers of subscribers, unless substantially more mobile-capable spectrum is allocated. Thus, the second mile bandwidth requirement of 1 to 5 Gbps of bandwidth can be scaled back to perhaps 1 Gbps, or even less, in the wireless last mile context.

II. TECHNOLOGY OPTIONS FOR PROVISION OF TRANSPORT

A. Middle Mile

15. In general, the middle mile bandwidth requirement of 5 to 62.5 Gbps can only be satisfied by fiber optic-based technology, particularly over extended distances greater than five kilometers. In particular, I believe fiber-based metro Ethernet networks will

play a primary role in the future for middle mile applications. With currently available fiber optic systems, the lower bound of 5 Gbps of middle mile capacity can most simply and cost-effectively be provided by a single 10 Gbps Ethernet (10GE) circuit over a single fiber pair. The corresponding upper bound of 62.5 Gbps of capacity can only be met using multiple 10GE circuits implemented over multiple fiber pairs or a combination of wave division multiplexing over a single fiber pair. Appropriate sized connection-oriented carrier Ethernet/MPLS switches with multiple 10GE and 1GE are now commercially available, and their costs are declining due to a combination of advances in technology and electronics vendor competition. Unamplified 10GE metro Ethernet circuits can reach up to 80 km over single mode fiber today, typically enough reach for urban and suburban applications. Rural applications with middle miles measured in the hundreds of kilometers require optical amplification at least every 80-100 km, adding to the one time electronics cost. In addition, 100 Gbps Ethernet ("100GE") capable packet switches will become commercially available in 2010, and the first 100GE capable metro transport systems will be deployed in the same time frame. Going forward, an aggregating Ethernet/MPLS switch complex in the central office/cable head end/MSC location will require only a small number of fiber pairs to drive hundreds of Gbps of middle mile capacity, and all of this technology will be commercially available in the next three years.

16. The cost of fiber-based middle mile links is proportional to their length. In urban or suburban areas, the cost of new fiber network construction varies widely, roughly from $4 to $35 per foot, exclusive of electronics costs and depending on whether the fiber

cable will be pole attached or buried, the number of fiber stands in the cable, and right-of-way costs. In general, in urban and suburban areas, the dominant cost for fiber deployment is the installation cost, not the fiber cable cost. Fiber construction costs in rural areas tends to be less expensive than in urban and suburban areas, since aerial fiber is often deployed and trenching costs can be lower.

17. While I believe that fiber-based metro Ethernet networks will become the primary means of meeting middle mile bandwidth requirements, it should be noted that the incumbent LECs have yet to an implement an Ethernet transport offering that is commercially attractive to XO and other wholesale customers. As an initial matter, AT&T and Verizon have not integrated the local access Ethernet services that were offered by AT&T and MCI at the time of their acquisition with their own Ethernet products. As a result, neither AT&T nor Verizon provides a single, uniform Ethernet service with the same technical specifications across its entire regional service area. Moreover, these incumbent LECs' first-generation Ethernet offerings suffer from a number of technical and operational limitations. Verizon and AT&T cannot guarantee adequate levels of network availability, latency, and throughput on their Ethernet networks, and the cost-per-bit of these services is well above what XO and other competitive LECs need to serve their customers. For these reasons, XO to date has not made significant use of the Ethernet offerings of AT&T and Verizon.

18. Besides fiber, point-to-point fixed wireless systems could address the low end of the bandwidth range required in the middle mile. For example, current radio systems are capable of 1 Gbps, and higher capacity systems are under development. In any case,

multiple parallel fixed wireless links would be necessary to implement the middle mile, and the distances would be limited to several kilometers. Longer fixed wireless links are possible, but they would require one or more regeneration sites and would not be cost competitive to fiber.

19. Going forward, current copper-based technologies such as DS1 and DS3 will play an important role for middle mile transport during a transitional period as new higher bandwidth networks are deployed. One mid-sized customer today can easily consume the entire 45 Mbps of bandwidth of a DS3. SONET-based circuits such as OC-3c (155 Mbps), OC-12c (622 Mbps), OC-48c (2.5 Gbps) and OC-192c (10 Gbps) are broadly deployed in Internet networks today, but only OC-192c or multiple OC-48c links have sufficient bandwidth to meet even the lower end of the 5 – 60+ Gbps of bandwidth required in the middle mile going forward. On a cost-per-bit basis, SONET links and router ports are not cost competitive with Ethernet-based links and routers, with SONET technology being anywhere from 2 to 5 times more expensive.

B. Second Mile

20. As described above, the bandwidth requirement for second mile broadband access is projected to range from 1 Gbps to 5 Gbps. This bandwidth requirement can be cost-effectively met today via fiber or fixed wireless technologies. A single fiber pair can today provide the required second mile bandwidth using a single 1GE or 10GE circuit. Similarly, one or several parallel fixed wireless links are currently capable of providing the required second mile bandwidth. XO is a large scale holder of LMDS licenses, and